



NAVIGATING THE AI REVOLUTION: OPPORTUNITIES AND RISKS IN CAPITAL MARKETS

Abstract

Artificial intelligence is rapidly reshaping capital markets across trading, compliance, client communications, and risk management, delivering clear efficiency and decision-making gains while also creating new governance, regulatory, and systemic risks. As adoption accelerates ahead of consistent oversight, institutions must balance innovation with stronger controls, accountability, and resilient operating frameworks.



What AI makes possible

The efficiency case has been proven to be robust. Industry research identifies three well-observed benefits: superior predictive performance over traditional statistical models, automation of repetitive back-office tasks, and improved risk management through better liquidity forecasting and anomaly detection.

Liquidity risk is a useful illustration. Traditional statistical methods failed to predict liquidity deterioration during the 2007–2009 financial crisis. AI models, by incorporating social media sentiment

and non-linear variable interactions, have demonstrated substantially better prediction under stressed economic conditions. Better liquidity forecasting affects position sizing, margin management, and regulatory capital. These are major considerations for a capital markets desk.

In AML, the improvement is operational. Rule-based approaches generate high volumes of false positives requiring costly human review. AI models are reducing false positive rates while continuously adapting

to new laundering patterns, something static rules cannot do.

For hedging, reinforcement learning models produced bottom-line improvements over delta-hedging baselines in backtested environments, though deployment in live markets remains complicated by the challenge of ensuring simulated data behaviour correlates with real conditions. Backtested performance and live performance diverge in practice.

The risk landscape

Some risks have already been documented to cause harm. While there was speculation around these, there was no inherent defence.

Market manipulation

Algorithmic market manipulation is among the most serious concerns. There has been credible research whose findings are difficult to dismiss: AI systems can collude to fix prices and sideline human traders

without explicit programming instructions. Even simple algorithms can develop collusion-like behaviour autonomously, in certain conditions. Critics dispute the scale of the effect. Some argue that AI's impact on market efficiency is ambiguous rather than clearly negative. But the core finding is that AI can generate market conduct outcomes that no human intended or sanctioned.

Herding

Herding is a related but distinct risk. When many institutions run models trained on similar datasets, those models respond similarly to the same market signal. At scale, that correlation amplifies markets. A moderate sell-off can become a liquidity cascade when all the models read the same conditions the same way at the same moment.

Fraud

AI-enhanced fraud already shows measurable risk. Research suggests that investors exposed to AI-enhanced scams committed 22% more capital than those who encountered conventional fraud. Deepfakes, AI-generated phishing, and identity manipulation are documented threats. The SEC has initiated enforcement

actions against firms that made misleading AI capability claims in disclosure documents. This has been termed "AI washing". The first investor lawsuit against an AI developer for trading losses reportedly occurred as early as 2019.

At the model level, additional vulnerabilities apply. AI systems can hallucinate and produce plausible but

factually wrong outputs. A hallucinated data point in a risk model has a more severe impact than a rounding error. Models drift over time as real market conditions diverge from training data. There has also been strong criticism over the new concern of model collapse. Model collapse is where AI systems trained on AI-generated data degrade progressively.

Concentration and the systemic question

Experts have identified third-party concentration as a structural risk that may be more consequential than any individual AI failure. Developing state-of-the-art models requires computing resources and capital that only a small number of technology companies can sustain. A recent survey found that 40% of machine learning models at capital markets

institutions are implemented through vendor tools and cloud services. When that many institutions share the same providers, a single failure is a market-wide event.

The problem extends beyond primary vendors. The "nth-party risks", that are the chains of subcontractors and downstream dependencies that sit behind the primary AI vendor, are largely invisible

to the institutions that rely on them. Full supply chain visibility is operationally impossible in most cases. Cloud outages at major providers have already caused simultaneous disruptions across financial platforms. Some argue that professional management of large-scale infrastructure brings resilience; others see correlated exposure at the systemic scale.



Regulatory responses and different organisational approaches

Three distinct approaches have been documented. The first applies existing technology-neutral financial regulation to AI without modification. The second issues specific AI guidance within established frameworks. The third builds bespoke AI legislation.

Examples of regulatory responses are:

- The EU AI Act, which came into effect in August 2024, is the most ambitious example of the third. It classifies AI systems into risk tiers and imposes binding requirements on high-risk applications, including creditworthiness

evaluation.

- ESMA issued MiFID II-specific guidance on AI in retail investment services in May 2024.
- Canada's CSA published principles-based guidance in December 2024 covering governance, explainability, and transparency.
- Hong Kong's SFC issued a specific circular on generative AI in November 2024.

The picture in the US is less settled. There is no comprehensive federal AI legislation. The SEC launched an AI task force in

August 2025 and held a dedicated AI roundtable in May 2025, but a proposed rulemaking that had addressed predictive data analytics conflicts of interest was withdrawn in June 2025.

What is missing is convergence. The benefits and risks of AI depend heavily on the specific type of technology, how it was developed, and for what purpose, since not all AI is alike. Treating AI as a single regulatory category risks over-regulating low-risk applications while leaving genuinely dangerous ones under-examined.





Towards a responsible AI framework

Six measures for responsible AI deployment have emerged as the international reference standard:

- Senior management oversight
- Testing and monitoring
- Skills and expertise
- Third-party management
- Disclosure
- Data quality controls

Some frameworks have operationalised governance across seven trust attributes:

- Explainability
- Accountability
- Reliability
- Privacy
- Fairness
- Transparency
- Sustainability

This framework's core premise is that effective AI governance is not a compliance exercise; it is a continuous operational discipline. It must monitor for model

drift, test for bias, validate against real market conditions, and maintain clear accountability when things go wrong.

Human oversight remains non-negotiable, but the question is becoming harder to answer clearly. As agentic AI systems act at millisecond speeds with reduced human intervention, what meaningful oversight actually means in practice is still being defined. Regulators and institutions will need to answer this together before the technology outpaces the discussion.

End note

Artificial intelligence in capital markets is simultaneously the most powerful efficiency lever available to financial institutions today, and one of the most complex governance challenges the industry has confronted. The same capability that reduces AML false positives and sharpens liquidity forecasting can, if ungoverned, enable algorithmic market

manipulation, amplify systemic correlation, and concentrate operational dependency in a handful of providers.

The regulatory environment will continue to develop. Implementation of the EU AI Act and emerging US legislative activity will each add definition. But institutions that wait for regulatory clarity before building governance frameworks may

be falling behind competitively: 49% of the market has already adopted AI, and governance frameworks are still catching up.

Enterprises need to start building [AI-enabled operating models](#) that balance the efficiency gains of automation with governance rigour, the kind regulators and investors are increasingly demanding.

For more information, contact infosysbpm@infosys.com



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.